

# MapReduce-based fuzzy c-means clustering algorithm: implementation and scalability

Simone A. Ludwig<sup>1</sup>

Received: 30 October 2014 / Accepted: 20 April 2015 / Published online: 29 April 2015  
© Springer-Verlag Berlin Heidelberg 2015

**Abstract** The management and analysis of big data has been identified as one of the most important emerging needs in recent years. This is because of the sheer volume and increasing complexity of data being created or collected. Current clustering algorithms can not handle big data, and therefore, scalable solutions are necessary. Since fuzzy clustering algorithms have shown to outperform hard clustering approaches in terms of accuracy, this paper investigates the parallelization and scalability of a common and effective fuzzy clustering algorithm named fuzzy c-means (FCM) algorithm. The algorithm is parallelized using the MapReduce paradigm outlining how the *Map* and *Reduce* primitives are implemented. A validity analysis is conducted in order to show that the implementation works correctly achieving competitive purity results compared to state-of-the-art clustering algorithms. Furthermore, a scalability analysis is conducted to demonstrate the performance of the parallel FCM implementation with increasing number of computing nodes used.

**Keywords** MapReduce · Hadoop · Scalability

## 1 Introduction

Managing scientific data has been identified as one of the most important emerging needs of the scientific community in recent years. This is because of the sheer volume and increasing complexity of data being created or collected, in

particular, in the growing field of computational science where increases in computer performance allow ever more realistic simulations and the potential to automatically explore large parameter spaces. As noted by Bell et al. [1]: “As simulations and experiments yield ever more data, a fourth paradigm is emerging, consisting of the techniques and technologies needed to perform data intensive science”. The question to address is how to effectively generate, manage and analyze the data and the resulting information. The solution requires a comprehensive, end-to-end approach that encompasses all stages from the initial data acquisition to its final analysis.

Data mining is a relatively broad field that deals with the automatic knowledge discovery from databases, and is one of the most developed fields in the area of artificial intelligence. Given the rapid growth of data collected in various fields and their potential usefulness requires efficient tools to extract and make use of the potentially gathered knowledge [2].

One of the important data mining tasks is classification, which is an effective method that is used in many different fields. The main idea behind the classification task is to build a model (classifier) that assigns items in a collection to target classes with the goal to accurately predict the target class for each item in the data [3]. There are many techniques that can be used to do a classification process such as decision trees, Bayes networks, genetic algorithms, genetic programming and many others [4]. Another important data mining technique used when analyzing data is clustering [5]. The main goal of clustering algorithms is to divide a set of unlabeled data objects into different groups called clusters (each group has common specifications between the group members). The cluster membership measure is based on a similarity measure. To obtain high quality clusters, the similarity measure between the data

---

✉ Simone A. Ludwig  
simone.ludwig@ndsu.edu

<sup>1</sup> Department of Computer Science, North Dakota State University, Fargo, ND, USA

objects in the same cluster is to be maximized, and the similarity measure between the data objects from different groups is to be minimized [6].

Clustering is the classification of objects into different groups, i.e., the partitioning of data into subsets (clusters), such that data in each subset shares some common features, often proximity according to some defined distance measure. Unlike conventional statistical methods, most clustering algorithms do not rely on the statistical distribution of data, and thus can be usefully applied in situations where little prior knowledge exists [7].

Most sequential classification/clustering algorithms suffer from the problem that they do not scale with larger sizes of data sets, and most of them are computationally expensive, both in terms of time and space. For these reasons, the parallelization of the data classification/clustering algorithms is paramount in order to deal with large scale data. To develop a good parallel classification/clustering algorithm that takes big data into consideration, the algorithm should be efficient, scalable and obtain high accuracy solutions.

In order to enable big data to be processed, the parallelization of data mining algorithms is paramount. Parallelization is a process where the computation is broken up into parallel tasks. The work done by each task, often called its grain size, can be as small as a single iteration in a parallel loop or as large as an entire procedure. When an application can be broken up into large parallel tasks, the application is called a coarse grain parallel application. Two common ways to partition computation are task partitioning, in which each task executes a certain function, and data partitioning, in which all tasks execute the same function but on different data.

This paper proposes the parallelization of a fuzzy c-means (FCM) clustering algorithm. The parallelization methodology used is the divide-and-conquer methodology referred to as MapReduce. The implementation details are explained in details outlining how the FCM algorithm can be parallelized. Furthermore, a validity analysis is conducted in order to demonstrate the correct functioning of the implementation measuring the purity and comparing these to state-of-the art clustering algorithms. Moreover, a scalability analysis is conducted to investigate the performance of the parallel FCM implementation by measuring the speedup for increasing number of computing nodes used.

The remainder of this paper is as follows. Section 2 introduces clustering and fuzzy clustering in particular. The following section (Sect. 3) discusses related work in the area of big data processing. In Sect. 4, the implementation is described in details. The experimental setup and results are given in Sects. 5 and 6 concludes this work outlining the findings obtained.

## 2 Background to clustering

Clustering can be applied to data that are quantitative (numerical), qualitative (categorical), or a mixture of both. The data usually are observations of some physical process. Each observation consists of  $n$  measured variables (features), grouped into an  $n$ -dimensional column vector  $z_k = [z_{1k}, \dots, z_{nk}]^T$ ,  $z_k \in \mathbb{R}^n$ .

A set of  $N$  observations is denoted by  $Z = \{z_k | k = 1, 2, \dots, N\}$ , and is represented as an  $n \times N$  matrix:

$$Z = \begin{pmatrix} z_{11} & z_{12} & \dots & z_{1N} \\ z_{21} & z_{22} & \dots & z_{2N} \\ \dots & \dots & \dots & \dots \\ z_{n1} & z_{n2} & \dots & z_{nN} \end{pmatrix}. \quad (1)$$

There are several definitions of how a cluster can be formulated depending on the objective of clustering. In general, a cluster is a group of objects that are more similar to one another than to members of other clusters [8, 9]. The term “similarity” is defined as mathematical similarity and can be defined by a distance norm. Distance can be measured among the data vectors themselves or as a distance from a data vector to some prototypical object (prototype) of the cluster. Since the prototypes are usually not known beforehand, they are determined by the clustering algorithms simultaneously with the partitioning of the data. The prototypes can be vectors of the same dimension as the data objects, however, they can also be defined as “higher-level” geometrical objects such as linear or nonlinear subspaces or functions. The performance of most clustering algorithms is influenced by the geometrical shapes and densities of the individual clusters, but also by spatial relations and distances among the clusters. Clusters can be well-separated, continuously connected or overlapping [7].

Many clustering algorithms have been introduced and clustering techniques can be categorized depending on whether the subsets of the resulting classification are *fuzzy* or *crisp* (hard). *Hard clustering* methods are based on classical set theory and require that an object either does or does not belong to a cluster. Hard clustering means that the data is partitioned into a specified number of mutually exclusive subsets. *Fuzzy clustering* methods, however, allow the objects to belong to several clusters simultaneously with different degrees of membership. Fuzzy clustering is seen as more natural than hard clustering since the objects on the boundaries between several classes are not forced to fully belong to one of the classes, but rather are assigned membership degrees between 0 and 1 indicating their partial membership. The concept of *fuzzy partition* is vital for cluster analysis, and therefore, also for the identification techniques that are based on

fuzzy clustering. Fuzzy and possibilistic partitions are seen as a generalization of *hard partition* which is formulated in terms of classical subsets [7].

The objective of clustering is to partition the data set  $Z$  into  $c$  clusters. For now let us assume that  $c$  is known based on prior knowledge. Using classical sets, a *hard partition* of  $Z$  can be defined as a family of subsets  $A_i$   $1 \leq i \leq c \subset P(Z)^1$  with the following properties [7, 8]:

$$\bigcup_{i=1}^c A_i = Z \tag{2a}$$

$$A_i \cap A_j = \emptyset, \quad 1 \leq i \neq j \leq c \tag{2b}$$

$$\emptyset \subset A_i \subset Z, \quad 1 \leq i \leq c. \tag{2c}$$

Equation 2a denotes that the union subset  $A_i$  contains all the data. The subsets have to be disjoint as stated by Eq. 2b, and none of them can be empty nor contain all the data in  $Z$  as given by Eq. 2c. In terms of *membership (characteristic) function*, a partition can be conveniently represented by the *partition matrix*  $U = [\mu_{ik}]_{c \times N}$ . The  $i$ th subset  $A_i$  of  $Z$ . It follows from Eqs. 2 that the elements of  $U$  must satisfy the following conditions [7]:

$$\mu_{ik} \in \{0, 1\}, \quad 1 \leq i \leq c, \quad 1 \leq k \leq N \tag{2d}$$

$$\sum_{i=1}^c \mu_{ik} = 1, \quad 1 \leq k \leq N \tag{2e}$$

$$0 < \sum_{k=1}^N \mu_{ik} < N, \quad 1 \leq i \leq c. \tag{2f}$$

Thus, the space of all possible hard partition matrices for  $Z$  referred to as the hard partitioning space is defined by [7]:

$$M_{HC} = \left\{ U \in \mathbb{R}^{c \times N} \mid \mu_{ik} \in \{0, 1\}, \quad \forall i, k; \sum_{i=1}^c \mu_{ik} = 1, \right. \\ \left. \forall k; 0 < \sum_{k=1}^N \mu_{ik} < N, \quad \forall i \right\}. \tag{2g}$$

Generalization of the hard partitioning to the fuzzy partitioning follows by allowing  $\mu_{ik}$  to obtain values in  $[0, 1]$ . The conditions for a fuzzy partition matrix, analogous to hard clustering equations, is given by [7, 10]:

$$\mu_{ik} \in [0, 1], \quad 1 \leq i \leq c, \quad 1 \leq k \leq N \tag{3a}$$

$$\sum_{i=1}^c \mu_{ik} = 1, \quad 1 \leq k \leq N \tag{3b}$$

$$0 < \sum_{k=1}^N \mu_{ik} < N, \quad 1 \leq i \leq c. \tag{3c}$$

The  $i$ th row of the fuzzy partition matrix  $U$  contains values of the  $i$ th *membership function* of the fuzzy subset  $A_i$  of  $Z$ . Equation 3b constrains the sum of each column to 1, and thus the total membership of each  $z_k$  in  $Z$  equals one. The fuzzy partitioning space for  $Z$  is the set [7]:

$$M_{FC} = \left\{ U \in \mathbb{R}^{c \times N} \mid \mu_{ik} \in [0, 1], \quad \forall i, k; \sum_{i=1}^c \mu_{ik} = 1, \right. \\ \left. \forall k; 0 < \sum_{k=1}^N \mu_{ik} < N, \quad \forall i \right\}. \tag{3d}$$

### 3 Related work

In this section, first related work in the area of clustering is introduced, and afterwards clustering techniques that make use of parallelization mechanisms applied to big data are described. Given that this paper is concerned with the implementation and evaluation of a parallelized FCM algorithm, the related work outlines other research conducted in this area.

Fuzzy clustering can be categorized into three categories: hierarchical fuzzy clustering methods, graph-theoretic fuzzy clustering methods, and fuzzy clustering based on objective functions [11].

Hierarchical clustering techniques generate a hierarchy of partitions by means of agglomerative and divisive methods [11]. The agglomerative algorithms produce a sequence of clusters of decreasing number by merging two clusters from the previous level. The divisive algorithms perform the clustering the other way around [6]. Lee [12] proposed a hierarchical clustering algorithm in the area of business systems planning. The best cluster number is determined by a matching approach. Another technique called fuzzy equivalent relation-based hierarchical clustering handles the clustering without needing a predefined number of clusters [13].

Graph-theoretic fuzzy clustering methods are based on the notion of connectivity of nodes of a graph representing the data set. In graph-theoretic fuzzy clustering, the graph representing the data structure is a fuzzy graph and different types of connectivity lead to different types of clusters. The idea of fuzzy graphs was first mentioned in [14], whereby the fuzzy analogues of several basic graph-theoretic concepts such as bridges, cycles, paths, trees were introduced. Fuzzy graphs were first used for cluster analysis in [15].

Fuzzy clustering based on objective functions has shown to give the most precise formulation of the clustering. The FCM clustering model was first introduced in 1974 [16],

and was later extended and generalized in [8]. Since then, some variations of the method and model improvements were suggested.

The Gustafson–Kessel algorithm [17] is a fuzzy clustering technique that estimates the local covariance by partitioning the data into subsets that can be fitted with linear submodels. However, considering a general structure for the covariance matrix can have a substantial effect on the modeling approach, and therefore, the Gath–Geva algorithm [18] was proposed to overcome this. The fuzzy c-varieties [19] clustering algorithm is a fuzzy clustering algorithm where the prototype of each cluster is of type multi-dimensional linear.

By replacing the Euclidean distances with other distance measures and enriching the cluster prototypes with further parameters, other shapes than only spherical clusters can be discovered. Clusters might be ellipsoidal, linear, manifolds, quadrics or have even different volumes [11]. Fuzzy clustering has been proven to handle ambiguous data that share properties of different clusters using the idea of membership degrees to be assigned to data objects.

The use of fuzzy clustering, especially the FCM algorithm, has been successfully applied to image segmentation [20] where it showed to be very efficient. However, the FCM algorithm still lacks robustness to noise and outliers, especially in absence of prior knowledge of noise. A crucial parameter, used to balance between robustness to noise and effectiveness of preserving the details of the image, is manually set based on experience. In addition, the time of segmenting an image depends on the image size, and hence the larger the size of the image, the longer the segmentation time [21].

In [22], a k-means clustering variant is proposed. The approach is based on a prototype-based hybrid approach that speeds-up the k-means clustering method. The method first partitions the data set into small clusters (grouplets). Each grouplet is first represented by a prototype, and then the set of prototypes is partitioned into k clusters using the modified k-means method. The modified k-means method avoids empty clusters. Each prototype is replaced by its corresponding set of patterns to derive a partition of the data set. The proposed method has shown to be much faster in terms of processing time than basic k-means.

Big data clustering has recently received significant amount of attention. In particular, the aim is to build efficient and effective parallel clustering algorithms. Many of these algorithms use the MapReduce methodology that was proposed by Google [23]. In the following we will review research that has used and applied the MapReduce paradigm to the clustering of big data sets.

A MapReduce design and implementation of an efficient DBSCAN algorithm is introduced in [24]. The proposed algorithm addresses the drawbacks of existing parallel

DBSCAN algorithms such as the data balancing and scalability issues. The algorithm tackles these issues using a parallelized implementation that removes the sequential processing bottleneck and thereby improves the algorithm's scalability. Furthermore, the algorithm showed the largest improvement when the data was imbalanced. The evaluation conducted using large scale data sets demonstrate the efficiency and scalability of their algorithm.

A parallel k-means clustering algorithm based on MapReduce was proposed in [25]. The algorithm locates the centroids by calculating the weighted average of each individual cluster points via the *Map* function. The *Reduce* function then assigns a new centroid to each data point based on the distance calculations. At the end, a MapReduce iterative refinement technique is applied to locate the final centroids. The authors evaluated the implementation using the measures of speedup, scaleup, and sizeup. The results showed that the proposed algorithm is able to process large data sets of 8 GB using commodity hardware effectively.

In [26], an algorithm for solving the problem of document clustering using the MapReduce-based k-means algorithm is proposed. The algorithm uses the MapReduce paradigm to iterate over the document collection in order to calculate the term frequency and inverse document frequency. The algorithm represents the documents as (key, value) pairs, where the key is the document type and the value is the document text. The authors compare a non-parallelized version of k-means with the parallelized version to show the speedup gain. Furthermore, the experiments of the parallelized k-means algorithm on 50,000 documents showed that the algorithm performs very well in terms of accuracy for this text clustering task while having a reasonable execution time.

Another k-means clustering algorithm using MapReduce by merging the k-means algorithm with the ensemble learning bagging method is introduced in [27]. The proposed algorithm addresses the instability and sensitivity to outliers. Ensemble learning is a technique that uses a collection of models to achieve better results than any model in the collection, and bagging is one of the most popular type of ensemble techniques. The evaluation was performed on relatively small data sets (instances around 5000) and only four nodes were used for the parallelization. The authors show that a speedup is obtained on data sets consisting of outliers.

A self-organizing map (SOM) was modified to work with large scale data sets by implementing the algorithm using the MapReduce concept to improve the performance of clustering as shown in [28]. A SOM is an unsupervised neural network that projects high-dimensional data onto a low-dimensional grid and visually represents the

topological order of the original data. Unfortunately, the details of the MapReduce implementation are not given, but the experiments that were conducted with a small data set demonstrated the efficiency of the MapReduce-based SOM.

In [29], the authors applied the MapReduce framework to solve co-clustering problems by introducing a framework called distributed co-clustering with MapReduce. Unlike clustering which groups similar rows or columns independently, co-clustering searches for interrelated submatrices of rows and columns. The authors proved that using MapReduce is a good solution for co-clustering mining tasks by applying the algorithm to data sets such as collaborative filtering, text mining, etc. The experiments demonstrated that co-clustering with MapReduce can scale well with large data sets using up to 40 nodes.

In [30], a fast clustering algorithm with a constant factor approximation guarantee was proposed. The authors use a sampling technique to reduce the data size first, and then apply the Lloyd’s algorithm on the remaining data set. A comparison of this algorithm with several sequential and parallel algorithms for the k-median problem was conducted using randomly generated data sets to evaluate the performance of the algorithm. The randomly generated data sets contained up to 10 million points. The results showed that the algorithm achieves better or similar solutions compared to the existing algorithms especially on very large data sets.

A big data clustering method based on the MapReduce framework was proposed in [31]. The authors used an ant colony approach to decompose the big data into several data partitions to be used in parallel clustering. Applying MapReduce to the ant colony clustering algorithm lead to the automation of the semantic clustering to improve the data analysis task. The proposed algorithm was developed and tested on data sets with large number of records (up to 800 K) and showed acceptable accuracy with good speedup.

In [32], the authors introduced a new approach, called the best of both worlds method to minimize the I/O cost of cluster analysis with the MapReduce model by minimizing the network overhead among the processing nodes. They proposed a subspace clustering method to handle very large data sets in a reasonable amount of time. Experiments on terabyte data sets were conducted using 700 mappers and up to 140 reducers. The results showed very good speedup results.

As can be seen by the related work, different parallel clustering methods have been proposed in the past. The aim of this paper is to parallelize the FCM algorithm using the MapReduce concept in order to conduct a thorough experimentation and scalability analysis.

## 4 MapReduce-based fuzzy c-means implementation

### 4.1 Fuzzy c-means algorithm

Most analytical fuzzy clustering algorithms are based on the optimization of the basic c-means objective function, or some modification of the objective function. The optimization of the c-means functional represents a nonlinear minimization problem, which can be solved by using a variety of methods including iterative minimization [7]. The most popular method is to use the simple Picard iteration through the first-order conditions for stationary points, known as the FCM algorithm. Bezdek [8] has proven the convergence of the FCM algorithm. An optimal  $c$  partition is produced iteratively by minimizing the weighted within group sum of squared error objective function:

$$J = \sum_{i=1}^n \sum_{j=1}^c (u_{ij})^m d^2(y_i, c_j) \tag{4}$$

where  $Y = [y_1, y_2, \dots, y_n]$  is the data set in a  $d$ -dimensional vector space,  $n$  is the number of data items,  $c$  is the number of clusters which is defined by the user where  $2 \leq c \leq n$ ,  $u_{ij}$  is the degree of membership of  $y_i$  in the  $j$ th cluster,  $m$  is a weighted exponent on each fuzzy membership,  $c_j$  is the center of cluster  $j$ ,  $d^2(y_i, c_j)$  is a square distance measure between object  $y_i$  and cluster  $c_j$ . An optimal solution with  $c$  partitions can be obtained via an iterative process described in Algorithm 1. First, the fuzzy partition matrix is initialized, then within an iterative loop the following steps are executed: the cluster centers are updated using Eq. 5, then the membership matrix is updated using Eqs. 6 and 7. The loop is repeated until a certain threshold value is reached.

---

#### Algorithm 1 FCM Algorithm

---

**Input:**  $c$ : centroid matrix,  $m$ : weighted exponent of fuzzy membership,  $\epsilon$ : threshold value used as stopping criterion,  $Y = [y_1, y_2, \dots, y_n]$ : data

**Output:**  $c$ : updated centroid matrix

Randomly initialize the fuzzy partition matrix  $U = [u_{ij}^k]$

**repeat**

    Calculate the cluster centers with  $U^k$ :

$$c_j = \frac{\sum_{i=1}^n (u_{ij}^k)^m y_i}{\sum_{i=1}^n (u_{ij}^k)^m} \tag{5}$$

    Update the membership matrix  $U^{k+1}$  using:

$$u_{ij}^{k+1} = \frac{1}{\sum_{k=1}^c \frac{d_{ij}}{(d_{kj})^{\frac{2}{m-1}}}} \tag{6}$$

    where

$$d_{ij} = \|y_i - c_j\|^2 \tag{7}$$

**until**  $\max_{i,j} \|u_{ij}^k - u_{ij}^{k+1}\| < \epsilon$

Return  $c$

---



## 4.2 MapReduce-based fuzzy c-means algorithm (MR-FCM)

The growth of the internet has challenged researchers to develop new ideas to deal with the ever increasing amount of data. Parallelization of algorithms is needed in order to enable big data processing. Classic parallel applications that were developed in the past either used message passing runtimes such as message passing interface (MPI) [33] or parallel virtual machines [34].

A parallel implementation of the FCM algorithm with MPI was proposed in [35]. The implementation consists of three Master/Slave processes, whereby the first computes the centroids, the second computes the distances and updates the partition matrix as well as updates the new centroids, and the third calculates the validity index. Moderately sized data sets were used to evaluate the approach and good speedup results were achieved.

However, MPI utilizes a rich set of communication and synchronization constructs, which need to be explicitly programmed. In order to make the development of parallel applications easier, Google introduced a programming paradigm called MapReduce that uses the *Map* and *Reduce* primitives that are present in functional programming languages. The MapReduce implementation enables large computations to be divided into several independent *Map* functions. MapReduce provides fault tolerance since it has a mechanism that automatically re-executes *Map* or *Reduce* tasks that have failed.

The MapReduce model works as follows. The input of the computation is a set of key-value pairs, and the output is a set of output key-value pairs. The algorithm to be parallelized needs to be expressed by *Map* and *Reduce* functions. The *Map* function takes an input pair and returns a set of intermediate key-value pairs. The framework then groups all intermediate values associated with the same intermediate key and passes them to the *Reduce* function. The *Reduce* function uses the intermediate key and set of values for that key. These values are merged together to form a smaller set of values. The intermediate values are forwarded to the *Reduce* function via an iterator. More formally, the *Map* and *Reduce* functions have the following types:

$$\begin{aligned} \text{map}(k_1, v_1) &\rightarrow \text{list}(k_2, v_2) \\ \text{reduce}(k_2, \text{list}(v_2)) &\rightarrow \text{list}(v_3). \end{aligned}$$

In order to consider the mapping of the FCM algorithm to the *Map* and *Reduce* primitives, it is necessary for FCM to be partitioned into two MapReduce jobs since only one would not be sufficient. The first MapReduce job calculates the centroid matrix by iterating over the data records, and a second MapReduce job is necessary since the following

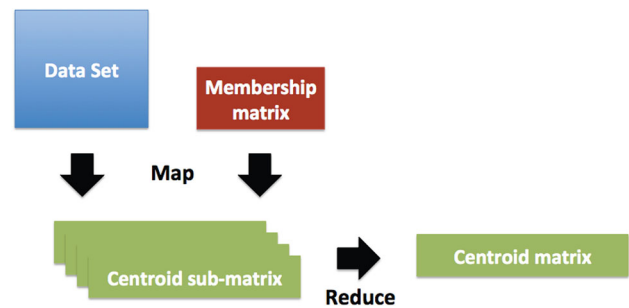


Fig. 1 First MapReduce job

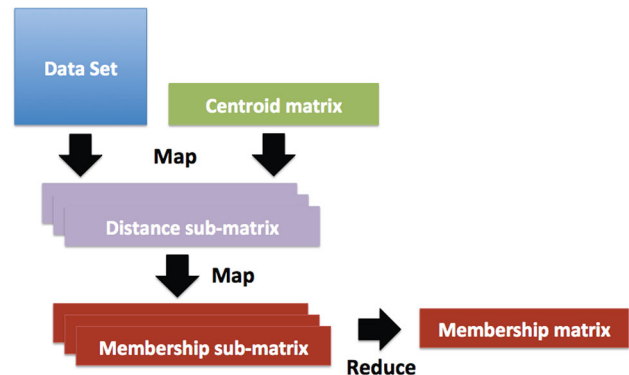


Fig. 2 Second MapReduce job

calculations need the complete centroid matrix as the input. The second MapReduce job also iterates over the data records and calculates the distances to be used to update the membership matrix as well as to calculate the fitness.

The details of the proposed implementation in the form of block diagrams are given in Figs. 1 and 2. As can be seen, in the first MapReduce job the mappers have a portion of the data set and a portion of the membership matrix and produce centroid sub-matrices. The reducer of the first MapReduce job then merges the sub-matrices into the centroid matrix.

The second MapReduce job compared to the first involves more computations to be executed. During the *Map* phase, portions of the data set are received and the distance sub-matrices and membership matrices, and sub-objective values are computed. Again, in the *Reduce* phase the membership sub-matrices are merged, and the sub-objective values are summed up. Please note that even though the figure shows three *Map* arrows, however, these steps are done in one *Map* phase.

An algorithmic description detailing the Main procedure and the two MapReduce jobs is given in Algorithms 2–6. The Algorithm 2 proceeds as follows. First, the membership matrix is randomly initialized. Then, the data set needs to be prepared such that the data set itself and the membership matrix are merged together vertically in order for

the first MapReduce job to have the data (both the data set as well as the membership matrix) available. This data set is stored in Hadoop distributed file system (HDFS) to be ready for the first MapReduce job. After this, the Hadoop framework calls the first and then the second MapReduce job. The resulting updated membership matrix is copied from the HDFS to the local file system for the second iteration to begin. The algorithm iterates calling the two mapreduce jobs several times until the stopping condition is met. Once this is completed the purity value is calculated using Eq. 8.

---

**Algorithm 2** Main Procedure of MR-FCM Algorithm
 

---

**Input:** data set

**Output:** purity

```

randomly initialize membership matrix
while stopping condition is not met do
    vertically merge data set with membership matrix and store in HDFS
    Hadoop calls map and reduce jobs of first and then second mapreduce operation
    copy membership matrix from HDFS and store locally
end while
calculate purity (Equation 8) and write result to file
  
```

---

Algorithm 3 shows the pseudo code for the *Map* function of the first MapReduce job. The inputs are the data record values and the membership matrix (which are vertically merged), and the output is the intermediate centroid matrix. During this *Map* function the intermediate centroid matrix values are calculated using Equation 5, which are then emitted.

---

**Algorithm 3** Map(key,value) of MapReduce job 1
 

---

**Input:** key: data record, value: data record values and membership matrix

**Output:** <key',value'> pair, where value' is the intermediate centroid matrix

```

for each key do
    calculate intermediate centroid matrix using Equation 5 and store in value'
    emit <key',value'> pair
end for
  
```

---

The intermediate centroid matrices are then merged by the *Reduce* function of the first MapReduce job 1 as shown in Algorithm 4.

---

**Algorithm 4** Reduce(key,value) of MapReduce job 1
 

---

**Input:** key: data record, value: intermediate centroid values

**Output:** <key',value'> pair, where value' is the centroid values

```

for each key do
    calculate centroid values by summing over intermediate centroid values and store in value'
    emit <key',value'> pair
end for
  
```

---

Algorithm 5 shows the *Map* function of the second MapReduce job. The function takes as the inputs the data

record values and the centroid matrix. The output is an intermediate membership matrix. This *Map* function first calculates the distances between the data points and the centroids using Equation 7, and then updates the intermediate membership matrix using Equation 6, which is then emitted.

---

**Algorithm 5** Map(key,value) of MapReduce job 2
 

---

**Input:** key: data record, value: data record values and centroid matrix

**Output:** <key',value'> pair, where value' is the intermediate membership matrix

```

for each key do
    calculate distances using Equation 7
    update intermediate membership matrix using Equation 6 and store in value'
    emit <key',value'> pair
end for
  
```

---

The *Reduce* function then merges the intermediate membership matrices and emits the membership matrix to be used in the next iteration. Algorithm 6 shows the details.

---

**Algorithm 6** Reduce(key,value) of MapReduce job 2
 

---

**Input:** key: data record, value: intermediate membership matrix

**Output:** <key',value'> pair, where value' is the membership matrix

```

for each key do
    merge intermediate membership matrices and store in value'
    emit <key',value'> pair
end for
  
```

---

## 5 Experiments and results

### 5.1 Clustering data set

The cover type data [36] set was used as the base for the experimentation. The data set contains data regarding 30 × 30 meter patches of forest such as elevation, distance to roads and water, hillshade, etc. The data set is used to identify which type of cover a particular patch of forest belongs to, and there are seven different types defined as output. The data set provides useful information to natural resource managers in order to allow them to make appropriate decisions about ecosystem management strategies. The data set characteristics are the following: number of instances is 581,012, number of attributes (categorical, numerical) is 54, and number of classes/labels is 7.

### 5.2 Computing environment

Most of the experiments were conducted on the Longhorn Hadoop cluster hosted by the Texas Advanced Computing Center (TACC)<sup>1</sup> consisting of 48 nodes containing 48 GB of RAM, 8 Intel Nehalem cores (2.5 GHz each), which

<sup>1</sup> <https://portal.longhorn.tacc.utexas.edu/>.

results in 384 compute cores with 2.304 TB of aggregated memory. For the experiments with the Mahout algorithms, another cluster, the Rustler Hadoop cluster<sup>2</sup> was used. The Rustler cluster consists of 66 nodes computing cluster for data intensive computing. Each node has dual eight core Ivy Bridge CPUs [Intel(R) Xeon(R) CPU E5-2650] running at 2.60 GHz, with 128 GB DDR3 memory and 16 1 TB SATA drives providing the backing for the HDFS file system. All nodes are connected via a 10 Gbps network. Hadoop version 0.21 was used for the MapReduce framework, and Java runtime 1.7 for the system implementation. The Apache Hadoop software library is a framework allowing distributed processing of large data sets across clusters of computers using simple programming models such as MapReduce [37]. Hadoop consists of MapReduce (processing element), and the HDFS (storage element).

### 5.3 Validation of the MR-FCM algorithm

In order to guarantee that the parallelized version of the FCM algorithm works appropriately, a performance measure needs to be used to quantify the resulting clustering quality. Given that the cover type data set is a “classification data set”, which includes the correct labels, therefore, the clustering quality can be measured in terms of purity [38]. Purity is defined as:

$$Purity = \frac{1}{n} \sum_{j=1}^k \max_i (|L_i \cap C_j|) \quad (8)$$

where  $C_j$  contains all data instances assigned to cluster  $j$  by the clustering algorithm,  $n$  is the number of data instances in the data set,  $k$  is the number of clusters that is generated from the clustering process,  $L_i$  denotes the true assignments of the data instances to cluster  $i$ .

In order to compare the MR-FCM algorithm, the following comparison algorithms are used: clustering toolkit (CLUTO) [39], approximate kernel possibilistic c-means (akPCM) [40], and approximate kernel fuzzy c-means (akFCM) [40]. The difference between the different algorithms are the membership update equations. CLUTO uses the hard c-means membership update, akPCM uses the possibilistic c-means membership update, and akFCM uses the FCM membership update, which our algorithm also uses. More details on the membership update equations can be found in [41]. In addition, the k-means as well as a recently released fuzzy k-means (FKM) algorithm (released in 2014) from the Mahout library [42] were also used for comparison.

The parameters for the proposed MR-FCM algorithm as well as for the FKM algorithm were set as follows:

**Table 1** Purity results for different clustering algorithms

Algorithm	Purity scores
CLUTO	0.49
k-means	0.50
akPCM	(RBF) 0.50 (D2) 0.49
akFCM	(RBF) 0.52 (D2) 0.53
FKM	0.52
MR-FCM	0.52

- cluster number = 7
- fuzzification exponent = 1.5
- epsilon = 1.0E−5

Table 1 lists the purity values for the different algorithms. For akPCM and akFCM the distinction is made using different kernels: Radial Basis Functions abbreviated as RBF, and degree-2 polynomial abbreviated as D2. We can see that our MR-FCM implementation achieves comparable purity results compared to akFCM and FKM, and better results than CLUTO, k-means and akPCM.

### 5.4 Scalability results

Since different data set sizes are necessary for the performance evaluation, the cover type data set has been split into different portions as shown in Table 2. Listed is the name of the data set, which corresponds to the number of rows contained, as well as the file size is given in bytes. Since the number of rows of the complete cover type data set is 581,012, for the experiments of the larger data set sizes, rows of the data sets were duplicated.

Three different experiments were conducted. The first experiment compares the ratio of the number of mappers and reducers used, and the second experiment uses the findings from the first experiment to perform a scalability analysis of the number of mappers/reducers used. The third experiments compares the Mahout FKM implementation with the proposed MR-FCM algorithm.

The difference between the two sets of measurements is that for the first experiment the same number of mappers and reducers are used for the second MapReduce job.

For the first MapReduce job preliminary experiments showed that the best performance is achieved when the number of reducers and the number of mappers is 7, since there are seven outcomes (clusters) in the data set. Therefore, in all experiments conducted the mappers and reducers for the first MapReduce job were set to 7, however, the timings of the first MapReduce (MR1) job are also be reported.

Figure 3 shows the execution time and the speedup of the second MapReduce (MR2) job. The left hand side (Fig. 3a, c) shows the results of using equal number of mappers

<sup>2</sup> <https://www.tacc.utexas.edu/systems/rustler>.



and reducers, and the right hand side (Fig. 3b, d) shows the results of using half the number of mappers for the reducers. The speedup results are calculated based on the time in ten mappers and ten reducers, and ten mappers and

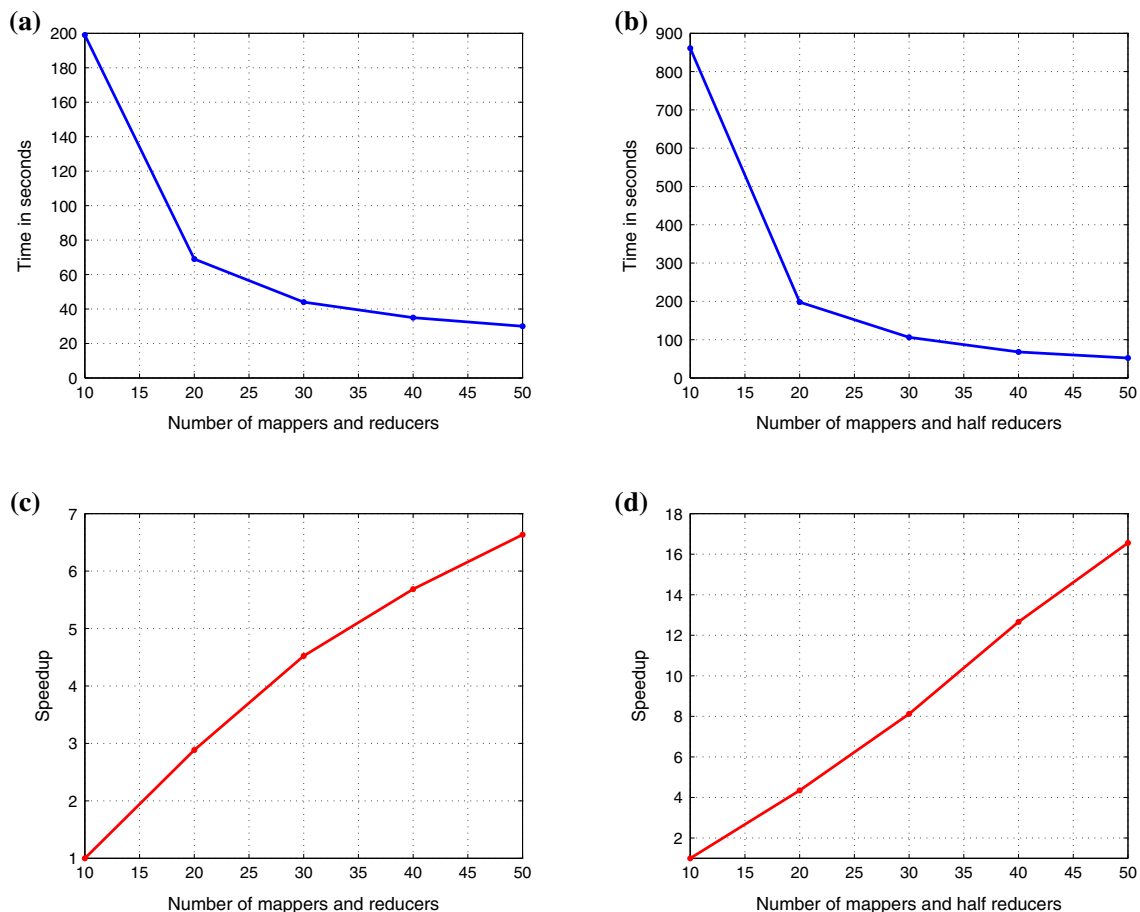
five reducers, as shown in Fig. 3c, d, respectively. What can be seen from this comparison is that the utilization of the Hadoop framework is much better when only half the number of mappers are used for the reducers. In particular, using ten mappers and ten reducers, the MR2 time is 861 s, whereas the MR2 time is 199 s when ten mappers and only five reducers are used. Therefore, an improvement of factor 4 is achieved with the setup of number of reducers equals half of the number of mappers. This finding is used throughout the next set of experiments.

The second experiment conducts a scalability analysis whereby the data set sizes are increased by 100,000 rows starting with a data set of 100,000 rows all the way up to 900,000 rows. Again, the number of mappers and reducers for the MR1 job was set to 7. Table 3 shows the execution time of MR1 for varying data set sizes. We observe a normal linear trend of the execution time for increasing data set sizes given that the number of mappers and reducers is fixed to 7.

Figure 4 shows the execution time of MR2 for increasing numbers of mapper and reducers used. The experiments use increments of 50, starting with 50 mappers

**Table 2** Data set description for performance evaluation

Data set	File size in bytes
100 K	18,237,388
200 K	36,606,436
300 K	55,011,262
400 K	73,387,535
500 K	91,761,302
600 K	109,998,690
700 K	128,367,738
800 K	146,772,564
900 K	165,148,837
1 M	183,522,604
2 M	367,045,208
3 M	550,567,812
4 M	734,090,416
5 M	917,613,020



**Fig. 3** Time in seconds and speedup results for MR2. **a** No. of mappers equals no. of reducers. **b** No. of reducers is half no. of mappers. **c** No. of mappers equals no. of reducers. **d** No. of reducers is half no. of mappers

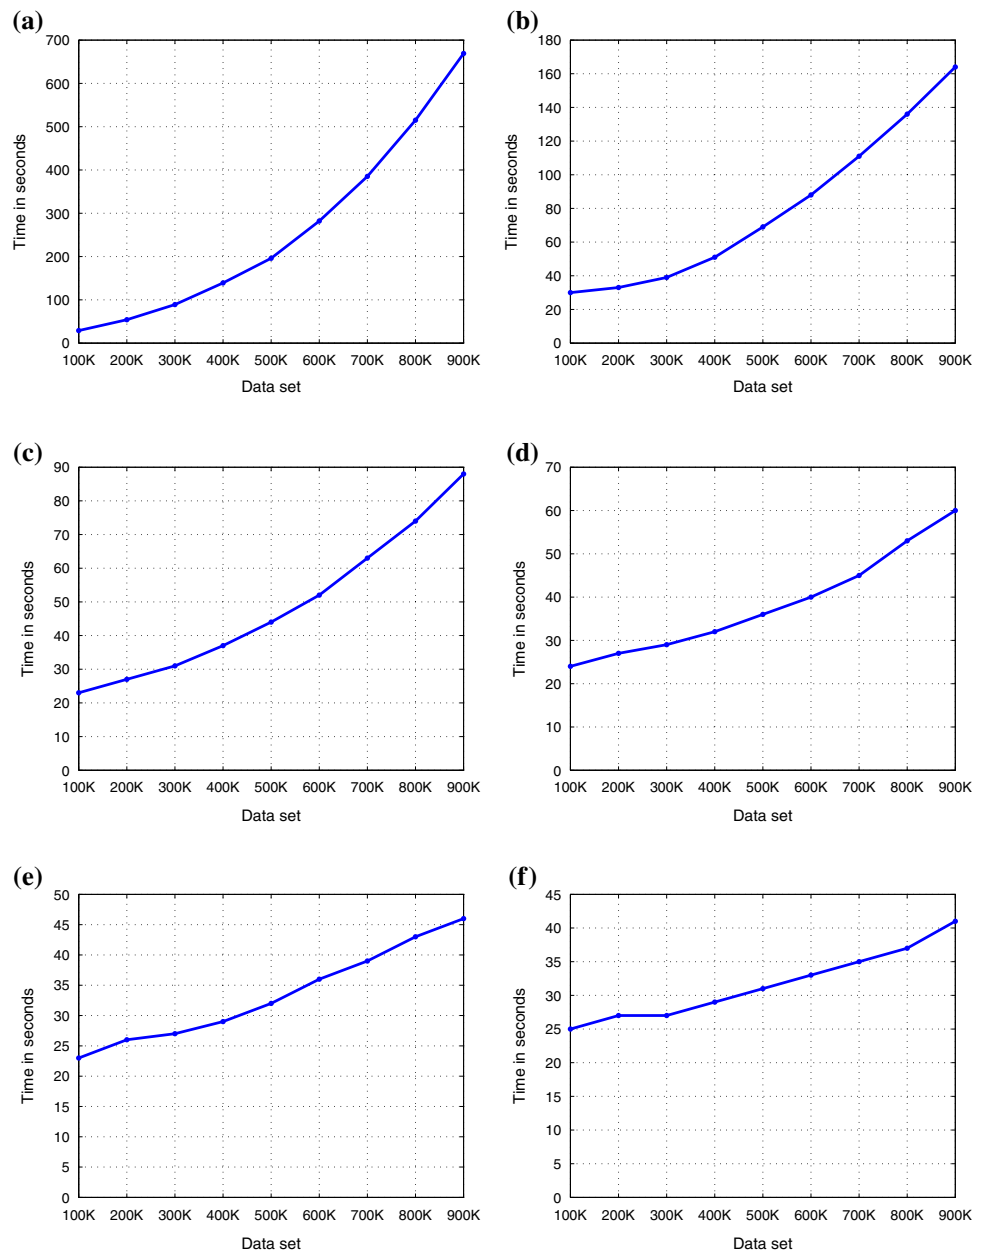
**Table 3** Time in seconds for MR1 for varying data set sizes (no. of mappers and reducers equals 7)

Data set (K)	MR1 time (s)
100	86 ± 3
200	156 ± 1
300	217 ± 2
400	285 ± 2
500	377 ± 11
600	429 ± 8
700	493 ± 9
800	553 ± 14
900	628 ± 9

(25 reducers) and ending with 500 mappers (250 reducers). Looking at all the figures we can see that the first figures with the lower numbers of mappers used show an “exponential” increase, whereas the later figures with the higher numbers of mappers used “flatten” and show an almost linear increase. The execution times for the 900 K data set are 669, 164, 88, 60, 46, and 41 s for 50, 100, 150, 200, 250, and 300 mappers, respectively.

The third experiment compares the two Mahout algorithms k-means and FKM with MR-FCM for data set sizes varying from 1 to 5 M. Since the infrastructure used at TACC does not allow the number of mappers and reducers

**Fig. 4** Time results of MR2 with varying data set sizes using different number of mappers and reducers. **a** 50 mappers and 25 reducers. **b** 100 mappers and 50 reducers. **c** 150 mappers and 75 reducers. **d** 200 mappers and 100 reducers. **e** 250 mappers and 125 reducers. **f** 300 mappers and 150 reducers



**Table 4** Time in seconds comparing k-means, FKM, and MR-FCM for varying data set sizes

Data set (M)	Time (k-means)	Time (FKM)	Time (MR-FCM)
1	352	769	797
2	365	773	925
3	368	789	1103
4	379	846	1318
5	388	875	1683

to be set explicitly, both algorithms are run without specifying the number of mappers/reducers in order to allow for a fair comparison. In addition, a maximum of 10 iterations are used to evaluate the execution time of both algorithms. Table 4 lists the results. We can see that k-means performs best as expected since it is the computationally less expensive algorithm. As for FKM and MR-FCM being very similar algorithms, we can see that the Mahout FKM implementation scales better than our MR-FCM algorithm. A possible reason for this might be that the Mahout library uses the vector format of the data set whereas the MR-FCM does not.

## 6 Conclusions

Since current clustering algorithms can not handle big data, there is a need for scalable solutions. Fuzzy clustering algorithms have shown to outperform hard clustering algorithms. Fuzzy clustering assigns membership degrees between 0 and 1 to the objects to indicate partial membership. This paper investigated the parallelization of the FCM algorithm and outlined how the algorithm can be parallelized using the MapReduce paradigm that was introduced by Google. Two MapReduce jobs are necessary for the parallelization since the calculation of the centroids need to be performed before the membership matrix can be calculated.

The accuracy of the MR-FCM algorithm was measured in terms of purity and compared to different clustering algorithms (both hard clustering and fuzzy clustering techniques) showed to produce comparable results.

The experimentation and scalability analysis revealed that the optimal utilization is achieved for the first MapReduce job using seven mappers and seven reducers, which is equal to the number of clusters in the data set. Furthermore, it was shown that for the second MapReduce job the best utilization is achieved when using half the number of mappers for the reducers. A factor of 4 in terms of speedup was achieved. The scalability analysis showed that for the data sets investigated (100 up to 900 K), a nearly linear increase can be observed when 250 mappers and 125 reducers and more are used. Another evaluation

compared the two Mahout algorithms k-means and FKM with MR-FCM for data set sizes varying from 1 to 5 M. This comparison showed that k-means, being the less computationally expensive algorithm performed best. FKM and MR-FCM are computationally very similar, however, the Mahout FKM algorithm showed to scale better than the MR-FCM algorithm.

Overall, with the implementation we have shown how a FCM algorithm can be parallelized using the MapReduce framework, and the experimental evaluation demonstrated that comparable purity results can be achieved. Furthermore, the MR-FCM algorithm scales well with increasing data set sizes as shown by the scalability analysis conducted.

Future work includes applying the MR-FCM algorithm to different clustering data sets emphasizing on the purity and scalability. Furthermore, larger data set sizes containing GBs of data should be investigated. For this however, another Hadoop cluster needs to be utilized where big data sets can be processed to achieve larger data clustering.

**Acknowledgments** The author acknowledges the Texas Advanced Computing Center (TACC) at the University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper.

## References

- Bell G, Hey AJG, Szalay A (2009) Beyond the data deluge. *Science* 323(AAAS):1297–1298. doi:10.1126/science.1170411
- Ghosh A, Jain LC (2005) Evolutionary computation in data mining series: studies in fuzziness and soft computing, vol 163. Springer, New York
- Tan P, Steinbach M, Kumar V (2005) Introduction to data mining. Addison-Wesley, New York. ISBN: 0-321-32136-7
- Jabeen H, Baig AR (2010) Review of classification using genetic programming. *Int J Eng Sci Technol* 2(2):94–103
- Han J (2005) Data mining: concepts and techniques. Morgan Kaufmann Publishers Inc., San Francisco
- Ludwig SA (2014) Clonal selection based fuzzy C-means algorithm for clustering. In: GECCO '14 Proceedings of the 2014 conference on genetic and evolutionary computation, pp 105–112
- Babuska R (2014) Fuzzy clustering lecture. <http://homes.di.unimi.it/~valenti/SlideCorsi/Bioinformatica05/Fuzzy-Clustering-lecture-Babuska>. Accessed Oct 2014
- Bezdek JC (1981) Pattern recognition with fuzzy objective function algorithms. Kluwer Academic Publishers, Norwell
- Jain AK, Dubes RC (1988) Algorithms for clustering data. Prentice-Hall Inc, Upper Saddle River
- Ruspini EH (1970) Numerical methods for fuzzy clustering. *Inf Sci* 2:319350
- Yang MS (1993) A survey of fuzzy clustering. *Math Comput Model* 18:1–16
- Lee HS (1999) Automatic clustering of business process in business systems planning. *Eur J Oper Res* 114:354–362
- Klir GJ, Yuan B (1995) Fuzzy sets and fuzzy logic theory and application. Prentice Hall PTR, Upper Saddle River
- Rosenfeld A (1975) Fuzzy graphs. In: Zadeh LA, Fu KS, Shimura M (eds) Fuzzy sets and their applications to cognitive and decision processes. Academic Press, New York

15. Matula DW (1970) Cluster analysis via graph theoretic techniques. In: Proceedings of the Louisiana conference on combinatorics, graph theory and computing, Winnipeg
16. Dunn JC (1973) A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters. *J Cybern* 3:32–57
17. Guerrero-Bote VP, Lopez-Pujalte C, de Moya-Anegon F, Herrero-Solana V (2003) Comparison of neural models for document clustering. *Int J Approx Reason* 34:287–305
18. Gath I, Geva AB (1989) Unsupervised optimal fuzzy clustering. *IEEE Trans Pattern Anal Mach Intell* 11(7):773–781
19. Bezdek JC, Coray C, Gunderson R, Watson J (1981) Detection and characterization of cluster substructure—linear structure, fuzzy c-varieties and convex combinations thereof. *SIAM J Appl Math* 40(2):358–372
20. Yang Y, Huang S (2007) Image segmentation by fuzzy c-means clustering algorithm with a novel penalty term. *Comput Inform* 26:17–31
21. Cai W, Chen S, Zhang D (2007) Fast and robust fuzzy c-means clustering algorithms incorporating local information for image segmentation. *Pattern Recognit* 40(3):825–838
22. Sarma TH, Viswanath P, Reddy BE (2013) A hybrid approach to speed-up the k-means clustering method. *Int J Mach Learn Cybern* 4:107–113
23. Dean J, Ghemawat S (2004) Mapreduce: simplified data processing on large clusters. In: Proceedings of the 6th conference on symposium on operating systems design and implementation (OSDI'04), vol 6, p 10
24. He Y, Tan H, Luo W, Feng S, Fan J (2014) Mr-dbscan: a scalable mapreduce-based dbscan algorithm for heavily skewed data. *Front Comput Sci* 8(1):83–99
25. Zhao W, Ma H, He Q (2009) Parallel k-means clustering based on mapreduce. In: Proceedings of the CloudCom'09. Springer, Berlin, pp 674–679
26. Zhou P, Lei J, Ye W (2011) Large-scale data sets clustering based on mapreduce and hadoop. *Comput Inf Syst* 7(16):5956–5963
27. Li H-G, Wu G-Q, Hu X-G, Zhang J, Li L, Wu X (2011) K-means clustering with bagging and mapreduce. In: Proceedings of the 44th Hawaii international conference on system sciences. IEEE Computer Society, Washington, DC, pp 1–8
28. Nair S, Mehta J (2011) Clustering with Apache Hadoop. In: Proceedings of the international conference, workshop on emerging trends in technology (ICWET'11), New York. ACM, New York, pp 505–509
29. Papadimitriou S, Sun J (2008) Disco: distributed co-clustering with map-reduce: a case study towards petabyte-scale end-to-end mining. In: Proceedings of the IEEE ICDM'08, Washington, DC, pp 512–521
30. Ene A, Im S, Moseley B (2011) Fast clustering using mapreduce. In: Proceedings of KDD'11. ACM, New York, pp 681–689
31. Yang J, Li X (2013) Mapreduce based method for big data semantic clustering. In: Proceedings of the 2013 IEEE international conference on systems, man, and cybernetics (SMC'13). IEEE Computer Society, Washington, DC, pp 2814–2819
32. Cordeiro F, Traina Jr C, Traina AJM, Lopez J, Kang U, Taloutos C (2011) Clustering very large multi-dimensional datasets with mapreduce. In: Proceedings of KDD'11. ACM, New York, pp 690–698
33. MPI (Message Passing Interface). <http://www-unix.mcs.anl.gov/mpi/>. Accessed 25 Apr 2015
34. PVM (Parallel Virtual Machine). <http://www.csm.ornl.gov/pvm/>. Accessed 25 Apr 2015
35. Modenesi MV, Costa MCA, Evsukoff AG, Ebecken NF (2007) Parallel fuzzy c-means cluster analysis. In: Lecture notes in computer science on high performance computing for computational science (VECPAR'06). Springer, New York
36. Blackard JA (1998) Comparison of neural networks and discriminant analysis in predicting forest cover types. Ph.D. dissertation, Department of Forest Sciences, Colorado State University, Fort Collins, Colorado
37. Apache Hadoop. <http://hadoop.apache.org/>. Accessed 25 Apr 2015
38. Han J (2005) Data mining: concepts and techniques. Morgan Kaufmann, San Francisco
39. Karypis G (2003) CLUTO: a clustering toolkit. University of Minnesota, Computer Science. Tech. Rep. 02-017
40. Havens TC, Chitta R, Jain AK, Rong J (2011) Speedup of fuzzy and possibilistic kernel c-means for large-scale clustering. In: Proceedings of IEEE international conference on fuzzy systems (FUZZ), pp 463–470
41. Hathaway R, Bezdek J (1995) Optimization of clustering criteria by reformulation. *IEEE Trans Fuzzy Syst* 3:241245
42. Mahout Library. <http://mahout.apache.org/>. Accessed 25 Apr 2015